ASAP | Machine Data Insights

Automated Splunk Admin Portal

The Automated Splunk Admin Portal (ASAP) is a Splunk® app that centralizes and automates Splunk administration tasks and provides configuration information, operational status and performance monitoring across *all* of your Splunk environments and servers from one web portal accessible on any PC, tablet, or smartphone

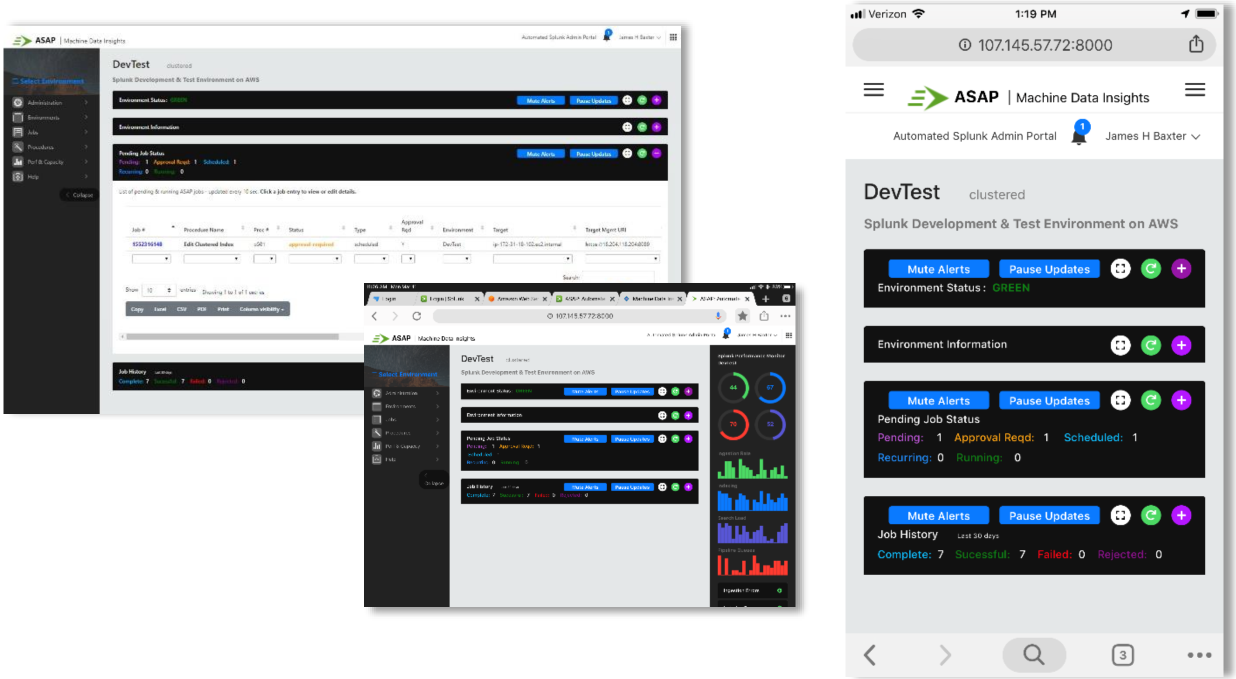
Verizon 1:19 PM

107.145.57.72:8000

ASAP | Machine Data Insights

Automated Splunk Admin Portal 1 James H Baxter

DevTest clustered

Splunk Development & Test Environment on AWS

Environment Status : GREEN

Environment Information

Pending Job Status

Pending: 1 Approval Reqd: 1 Scheduled: 1

Recurring: 0 Running: 0

Job History Last 30 days

Complete: 7 Sucessful: 7 Failed: 0 Rejected: 0

ASAP provides the following capabilities:

- Monitor status and performance of Splunk server health including cpu/mem/disk, ingestion, indexing, search.
- Put HW & SW complement and configuration information at your fingertips – exportable to CSV/Excel®
- View consolidated Splunk server messages from a selected environment with status and error alerting.
- Create & edit Splunk apps, authentication, indexes, inputs, outputs, users / roles / capabilities.
- Automate Splunk server configuration management with customized templates for each function / role.
- Automate formerly manual, multi-step, error-prone tasks such as distribution of indexes and apps in clustered environments via a small ASAP client app with API endpoints installed on cluster masters and deployers.
- Prepare Splunk admin tasks via the ASAP UI which creates Jobs that can be scheduled, executed immediately, and/or require approval by Sr. admins/management before implementation.
- Initiate jobs via REST API calls to ASAP from external CMDB solutions such as ServiceNow® to facilitate self-help and resource deployment (apps, servers, sensors, etc.) automation.
- Utilize ‘procedure’ files written in python to perform tasks – developing custom procedures is straightforward.
- Encourage and facilitate optimized, standardized configurations for higher consistency and reliability.
- Reduce Splunk admin workload to allow focus on higher ROI tasks - lowering resource requirements 20-60%.
- Monitor and manage your Splunk environments from any location 24x7 on your tablet or smartphone.

ASAP Features

Environment Status

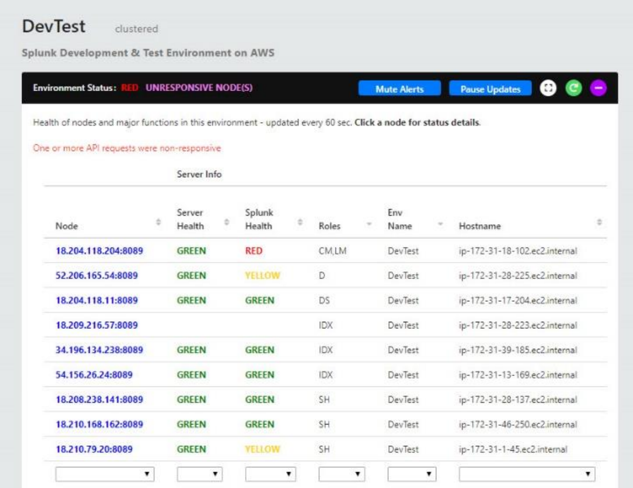

DevTest clustered

Splunk Development & Test Environment on AWS

Environment Status: RED UNRESPONSIVE NODE(S)

Health of nodes and major functions in this environment - updated every 60 sec. **Click a node for status details.**

One or more API requests were non-responsive

	Server Info				
Node	Server Health	Splunk Health	Roles	Env Name	Hostname
18.204.118.204:8089	GREEN	RED	CM,LM	DevTest	ip-172-31-18-102.ec2.internal
52.206.165.54:8089	GREEN	YELLOW	D	DevTest	ip-172-31-28-225.ec2.internal
18.204.118.11:8089	GREEN	GREEN	DS	DevTest	ip-172-31-17-204.ec2.internal
18.209.216.57:8089			IDX	DevTest	ip-172-31-28-223.ec2.internal
34.196.134.238:8089	GREEN	GREEN	IDX	DevTest	ip-172-31-39-185.ec2.internal
54.156.26.24:8089	GREEN	GREEN	IDX	DevTest	ip-172-31-13-169.ec2.internal
18.208.238.141:8089	GREEN	GREEN	SH	DevTest	ip-172-31-28-137.ec2.internal
18.210.168.162:8089	GREEN	GREEN	SH	DevTest	ip-172-31-46-250.ec2.internal
18.210.79.20:8089	GREEN	YELLOW	SH	DevTest	ip-172-31-1-45.ec2.internal

ASAP monitors the complement and status of Splunk server cpu/mem/disk components, as well as splunkd health metrics – Data Forwarding, File Monitor Inputs, Indexer Clustering, and Search Head connectivity.

Clicking or touching a table entry in the **Environment Status panel** drop-down opens a **Status Details** view with overall status, metrics, details, explanations, and related Splunk messages for reported anomalies.

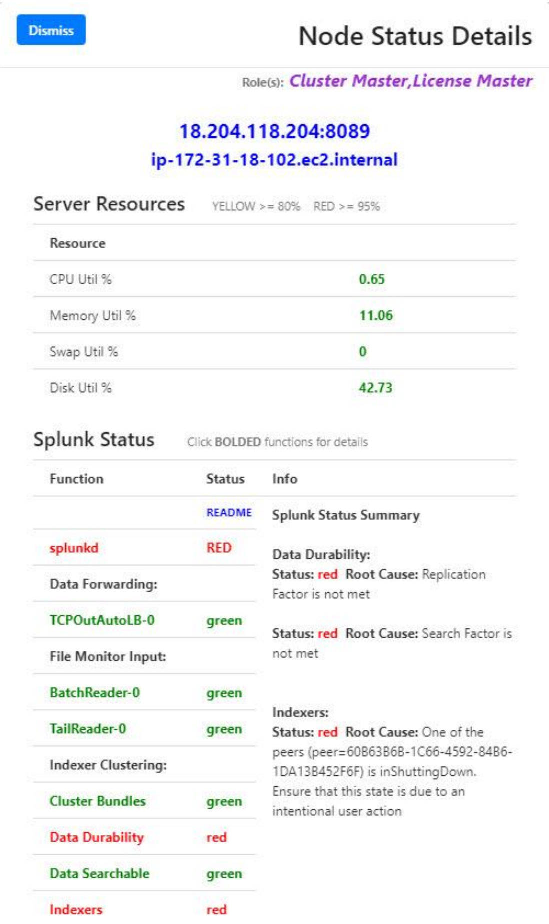

Node Status Details

Role(s): *Cluster Master,License Master*

18.204.118.204:8089
ip-172-31-18-102.ec2.internal

Server Resources YELLOW >= 80% RED >= 95%

Resource	
CPU Util %	0.65
Memory Util %	11.06
Swap Util %	0
Disk Util %	42.73

Splunk Status Click BOLDED functions for details

Function	Status	Info
	README	Splunk Status Summary
splunkd	RED	**Data Durability:** **Status:** red **Root Cause:** Replication Factor is not met
Data Forwarding:		
TCPOutAutoLB-0	green	**Status:** red **Root Cause:** Search Factor is not met
File Monitor Input:		
BatchReader-0	green	**Indexers:** **Status:** red **Root Cause:** One of the peers (peer=60B63B6B-1C66-4592-84B6-1DA13B452F6F) is inShuttingDown. Ensure that this state is due to an intentional user action
TailReader-0	green	
Indexer Clustering:		
Cluster Bundles	green	
Data Durability	red	
Data Searchable	green	
Indexers	red	

Environment Info

All the information you need about your Splunk servers (known as Nodes in ASAP) is available at your fingertips from the **Environment Info panel** for view or export – including:

- Hostname & Management URI
- Server roles & site – which node is **SH Captain**
- **Splunk GUID** – useful for troubleshooting
- CPU architecture & core counts
- OS, version, & build
- Splunk product type, version, & build
- Splunk License & state
- Forwarding state & KV Store status
- **Last time Splunk was started**
- **Transparent Huge Pages & ulimits metrics**

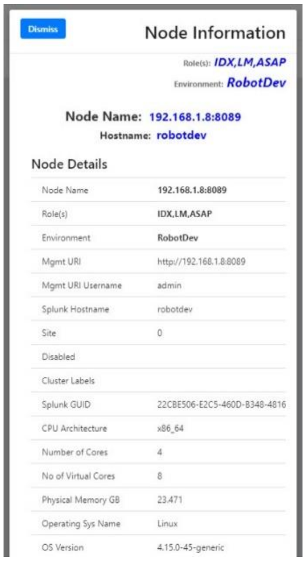

Node Information

Role(s): *IDX,LM,ASAP*
Environment: *RobotDev*

Node Name: 192.168.1.8:8089
Hostname: robotdev

Node Details

Node Name	**192.168.1.8:8089**
Role(s)	**IDX,LM,ASAP**
Environment	**RobotDev**
Mgmt URI	http://192.168.1.8:8089
Mgmt URI Username	admin
Splunk Hostname	robotdev
Site	0
Disabled	
Cluster Labels	
Splunk GUID	22CBE506-E2C5-460D-B348-4816
CPU Architecture	x86_64
Number of Cores	4
No of Virtual Cores	8
Physical Memory GB	23.471
Operating Sys Name	Linux
OS Version	4.15.0-45-generic

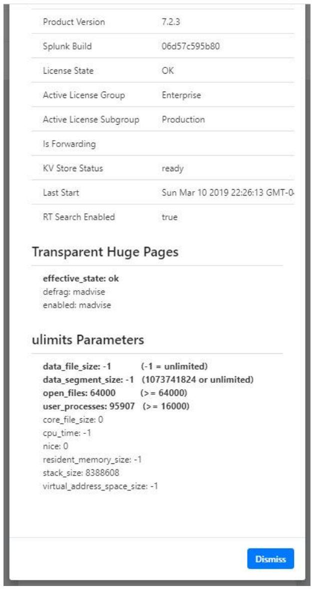

Product Version	7.2.3
Splunk Build	06d57c595b80
License State	OK
Active License Group	Enterprise
Active License Subgroup	Production
Is Forwarding	
KV Store Status	ready
Last Start	Sun Mar 10 2019 22:26:13 GMT-0
RT Search Enabled	true

Transparent Huge Pages

effective_state: ok
defrag: madvise
enabled: madvise

ulimits Parameters

```
data_file_size: -1          (-1 = unlimited)
data_segment_size: -1   (1073741824 or unlimited)
open_files: 64000         (>= 64000)
user_processes: 95907   (>= 16000)
core_file_size: 0
cpu_time: -1
nice: 0
resident_memory_size: -1
stack_size: 8388608
virtual_address_space_size: -1
```

ASAP Jobs

Submitting a Splunk admin task in ASAP results in the creation of a 'Job' which can be completed immediately or scheduled for off-hours execution.

Recurring jobs can also be created to periodically run custom check-and-resolve procedures such as checking indexer disk usage, knowledge bundle size, number of search artifacts on search heads, etc. and executing remedial / self-healing actions.

Job **Types** are set during creation:

immediate **scheduled** **recurring**

Job **Status** conditions are color coded:

new **approval-required** **approved**

rejected **running** **complete** **failed**

Jobs are executed by a multi-threaded 'jobrunner' which checks the job queue every 60 seconds.

Completed jobs with their status and run notes are moved to the Job History panel and viewable for (configurable) 1 – 30 days, as well as being logged in a job history file and immediately indexed for Splunk search retrieval.

Jobs can also be initiated via a REST API call to ASAP from external systems. These jobs can also be executed immediately, scheduled, or require approval as needed.

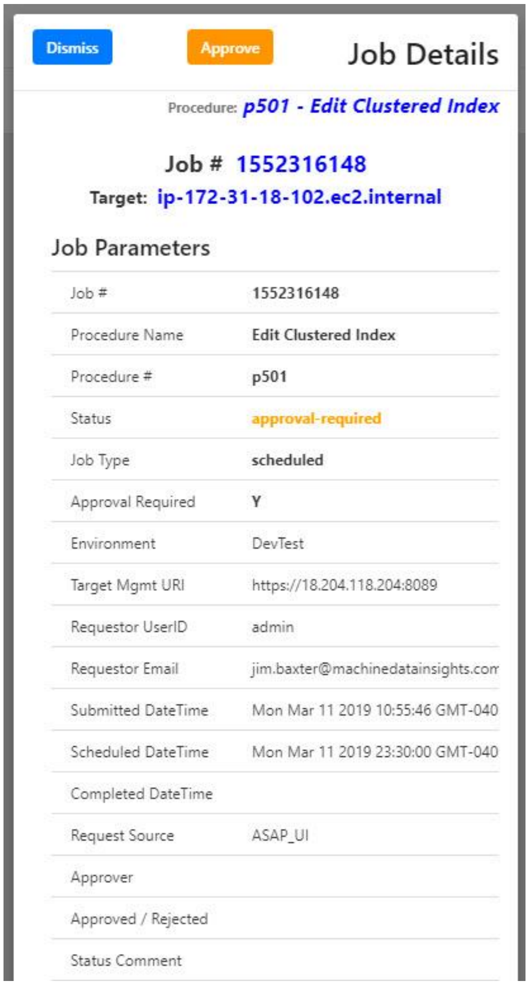

Job Details

Procedure: ***p501 - Edit Clustered Index***

Job # 1552316148

Target: ip-172-31-18-102.ec2.internal

Job Parameters

Job #	**1552316148**
Procedure Name	**Edit Clustered Index**
Procedure #	**p501**
Status	**approval-required**
Job Type	**scheduled**
Approval Required	**Y**
Environment	DevTest
Target Mgmt URI	https://18.204.118.204:8089
Requestor UserID	admin
Requestor Email	jim.baxter@machinedatainsights.com
Submitted DateTime	Mon Mar 11 2019 10:55:46 GMT-040
Scheduled DateTime	Mon Mar 11 2019 23:30:00 GMT-040
Completed DateTime	
Request Source	ASAP_UI
Approver	
Approved / Rejected	
Status Comment	

ASAP Procedures

- ASAP jobs run numbered 'procedures' which are python scripts that contain the logic and instructions for executing a Splunk admin task.
- Complex tasks such as creating an index in a clustered environment are orchestrated to configure the .conf file on the Cluster Master, validate the cluster bundle, apply the bundle to the indexing tier, and monitor the restart sequence to completion with status reporting.
- App deployments in distributed search head environments are accomplished in similar fashion from the Deployer.
- These typically manual and error-prone tasks are automated by procedures that leverage custom endpoints provided by a small ASAP client app installed on the Deployer and Cluster Master.
- Procedure actions are logged in Splunk at an INFO, ERROR, or configurable DEBUG level for process verification and troubleshooting.
- Along with the standard procedure scripts provided with ASAP, custom procedures can be quickly developed using templates and examples to address any unique administration requirement.

ASAP | Machine Data Insights

Administration with ASAP

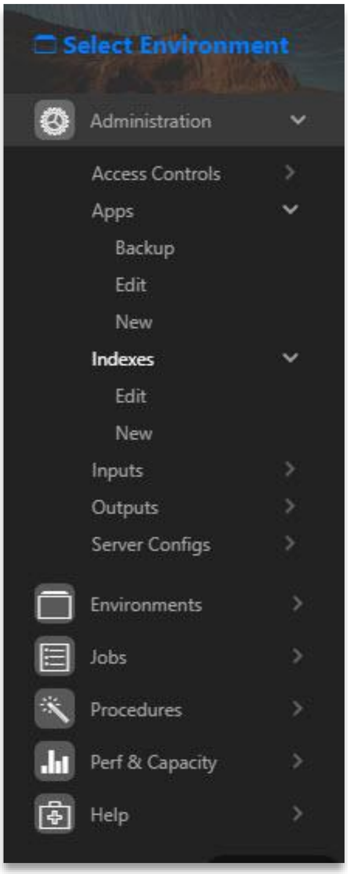

Selecting an Environment

Upon logging into ASAP you can click/touch **Select Environment** in the left-hand panel and choose the Splunk environment you want to work in.

You can also choose **All Nodes** to get a quick overview of the operational status of your entire Splunk infrastructure across all environments.

The left panel provides a menu of the most common Splunk admin activities. Clicking **Edit** opens a form with a table of existing entries, which can be selected and edited. Clicking **New** provides a form to provide configuration details and submit the job for execution.

Example: Create / Edit an Index

Creating or editing an index or any other Splunk entity is as easy as using Splunk Web but offers more functionality.

Note that in the **Create Index** form excerpt depicted here that you can set the Index Data Retention Period – an essential feature missing in Splunk Web which otherwise necessitates manual editing.

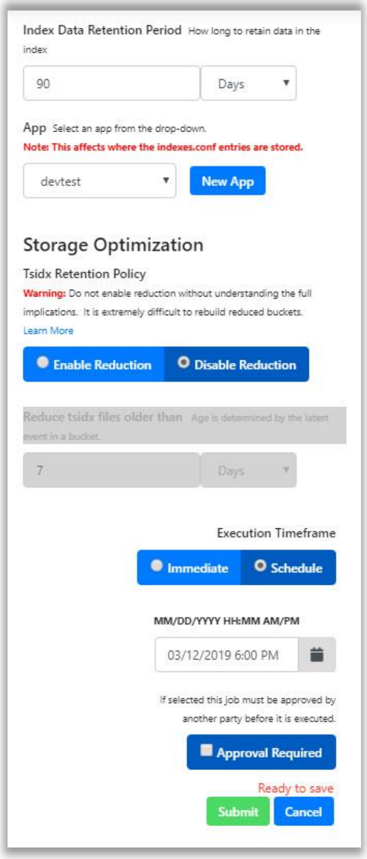

Index Data Retention Period How long to retain data in the index

90 | Days ▾

App Select an app from the drop-down.

Note: This affects where the indexes.conf entries are stored.

devtest ▾ | New App

Storage Optimization

Tsidx Retention Policy

Warning: Do not enable reduction without understanding the full implications. It is extremely difficult to rebuild reduced buckets.

Learn More

○ Enable Reduction ● Disable Reduction

Reduce tsidx files older than Age is determined by the latest event in a bucket.

7 | Days ▾

Execution Timeframe

○ Immediate ● Schedule

MM/DD/YYYY HH:MM AM/PM

03/12/2019 6:00 PM

If selected this job must be approved by another party before it is executed.

□ Approval Required

Ready to save

Submit | Cancel

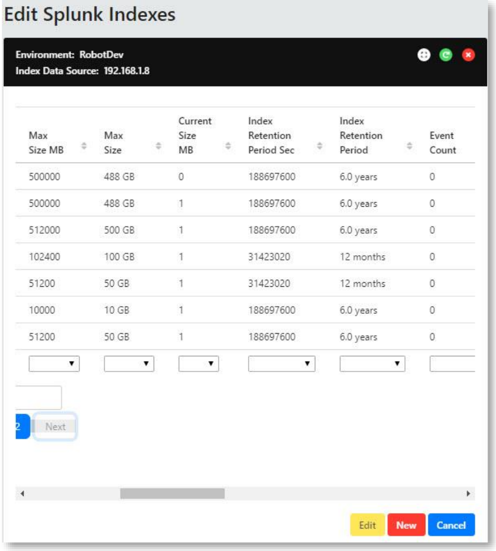

Edit Splunk Indexes

Environment: RobotDev
Index Data Source: 192.168.1.8

Max Size MB	Max Size	Current Size MB	Index Retention Period Sec	Index Retention Period	Event Count
500000	488 GB	0	188697600	6.0 years	0
500000	488 GB	1	188697600	6.0 years	0
512000	500 GB	1	188697600	6.0 years	0
102400	100 GB	1	31423020	12 months	0
51200	50 GB	1	31423020	12 months	0
10000	10 GB	1	188697600	6.0 years	0
51200	50 GB	1	188697600	6.0 years	0

This job has been scheduled for 6:00 PM to avoid execution during peak hours.

Clicking Submit after completing the form creates a **Job** which can be viewed in the Pending Job Status panel.

The **Edit Indexes** form provides a list of existing indexes and useful information about configured index size and data retention periods in summarized values, as well as type (events or metrics), app, event counts, earliest/latest events, and index storage paths.

ASAP Concepts

Environments and Nodes

ASAP functions within the context of Splunk **Environments** – such as the Dev/Test, QA/load test, and Production environments typical of most organizations, and **Nodes** which are the Splunk servers belonging to a specific environment. This allows automated tasks to be performed across the appropriate nodes in clustered environments without having to specify them in the ASAP UI.

Jobs and Procedures

Splunk administration is typically a hands-on function, with great opportunity for erroneous or sub-optimal configuration. ASAP procedures solidify processes and configurations, and jobs provide a change record.

ASAP Roles

ASAP leverages Splunk's native role and capability features to secure access to automation features with the addition of three ASAP-specific roles: an **asap-user** role can submit jobs which must be approved before execution by an **asap-power** role, which can execute any admin task, or an **asap-admin** role, which can also administer ASAP environments, nodes, and other internal ASAP settings.

Secure Portal to All Splunk Environments

All Splunk administration tasks conducted by ASAP are accomplished within the firewall rules already in place for a given environment. Cross-environment management by ASAP can be accomplished by opening a single API access port between ASAP nodes to provide multi-environment administration with minimal security concerns.

If company policy allows, the ASAP web interface can be accessed externally on admin and management PCs, tablets, and smartphones by opening a single port from the Internet to just the ASAP-hosting server. This allows the highest level of monitoring and management capabilities 24 x 7, again with minimal security concessions.

Why ASAP?

Most Splunk admin tasks are repetitive and tedious, and many require manual editing and deployment of complex configuration files, which is error prone and time consuming. Keeping configurations and processes consistent across an admin team of varying skill and experience levels is difficult. There is great need for an automation solution that provides configuration consistency and reliability while reducing admin workloads, as well as offering a Splunk configuration endpoint that interfaces well with external systems.

ASAP was developed to address this automation need, and to leverage the fact that admin tasks accomplished from Splunk Web or the Command Line Interface (CLI) actually creates an API call behind the scenes. This means that these API calls can be executed from any remote system, and that this remote system – ASAP – can monitor and manage any number of Splunk nodes from one portal, whereas Splunk Web / CLI are mostly local interfaces.

Learn More about ASAP and how it might fit in your organization, or schedule a demo by contacting:

Machine Data Insights

Jim.Baxter@MachineDataInsights.com
or visit:
https://www.machinedatainsights.com/asap.html